



APEX ENGLISH

Your voice. Global.

B2 · BUSINESS ENGLISH

UNIT 10 · B10A

AI at work

Theme · AI-augmented. Consulting mindset.



Unit 10 · B10A

AI at work

Theme: AI-augmented. Consulting mindset.



Learning objectives

- Talk honestly about where AI helps, where it hurts, and where it just sounds good in a slide.
- Use QUANTIFIERS (all, every, both, most, many, some, a few, few, hardly any, no / none) with precision — 'every' and 'most' are not interchangeable.
- Frame AI at work in terms of TRADE-OFFS, not hype: speed vs. quality, adoption vs. discernment, drafting vs. deciding.

Ask the teacher — notes to bring to class



Class Book

Warm-up · 'It drafts. Who decides?'

Think of **ONE** task you have handed to an AI tool in the last month at work or study. In 45 seconds out loud: (1) what did it draft, (2) how much of it survived your review, (3) what would have happened if you'd sent the first version untouched? End with 'The thing I now check every time is...'.
Was the draft 80% right and 20% quietly wrong — or the opposite?
Did you catch the errors, or did a colleague / client catch them?
Would you defend that draft in a meeting if someone challenged a number?
Bonus: which colleague on your team is over-trusting the tool right now?

Lead-in: name the quantifier trap

Match each real-sounding line (1–6) about AI at work to the quantifier trap it hides. Only **ONE** label per item.

1. 1. 'EVERY junior consultant is already using the tool — we're just formalising it.'

2. 2. 'MOST of the pushback is from a FEW people who don't understand AI.'

3. 3. 'ALL our client-facing decks now go through the copilot first.'

4. 4. 'BOTH legal AND risk have signed off.'

5. 5. 'NO consultant would send an AI draft without checking.'

6. 6. 'A FEW clients have asked whether we use AI on their work.'

followUp: Which of the six lines have YOU heard (or said) in the last quarter — and what did the speaker actually mean?

UNIVERSAL claim

CONTRADICTION inside one sentence

FALSE TOTALITY

PAIRED-UP SHIELD

ABSOLUTE DENIAL

HEDGED MINORITY

Reading: 'It drafts. We decide.'

Ask the teacher — notes to bring to class



For roughly two years, every professional-services firm in London, New York and Singapore has been quietly running the same experiment. Give a large-language-model copilot to a ring of consultants, watch the median deck-drafting time fall by something between thirty and fifty per cent, and then wait to see which firm has the first embarrassing public error. Most firms are now past that first error. The interesting question is no longer 'does the tool save time' — it obviously does — but 'which firms have built the discipline to keep the time savings without losing the client'.

The pattern in the firms that have handled it well is boringly consistent. All copilot output that touches a client — every memo, every slide, every model — is reviewed line-by-line by a named consultant at manager grade or above, who signs a shared log with their initials and the date. Not a review meeting; a signed line in a log. The rule is short enough to fit on a lanyard: it drafts, we decide. The firms that have handled it badly, by contrast, tend to share two habits. First, they conflate ADOPTION with DISCERNMENT — 'every consultant is using the tool' becomes evidence of success, when it is really just evidence of a login. Second, they measure only speed, and only in the direction that flatters them — median hours down, but never client-facing corrections up.

Both types of firm have, by now, discovered a strange second-order effect. The copilot is often best not at the task the partners bought it for (drafting) but at the task nobody quite planned for (compressing a 40-page transcript into a five-bullet 'here is what the client actually cares about'). A few firms have restructured whole workflows around this: the tool does the summary, a human writes the recommendation, and the two are stapled together with a signed cover page. Most firms are still awkwardly stuck in the middle — the tool drafts, the human 'sort of' reviews, and hardly anyone can tell you where the decision was made or by whom.

The final trap is the one nobody talks about in the town hall. If the tool is genuinely making your junior consultants two or three times more productive, the firm's answer to 'so what happens to the pyramid we've been billing against for forty years?' cannot be silence. Some firms are quietly widening the base — hiring fewer juniors, at higher standard, and telling them so. A few are widening the top — using the freed hours to send juniors into client meetings earlier. Almost none are keeping the pyramid exactly as it was and hoping nobody notices. The ones that are, are the ones whose best juniors leave first.

1. The interesting question about AI copilots in consulting is now...

- (A) Whether the tool saves time
- (B) Which firms build the discipline to keep the time savings without losing the client
- (C) Which model to buy

2. The 'It drafts. We decide.' discipline requires...

- (A) A weekly review meeting
- (B) Line-by-line review by a named manager-or-above who signs a shared log
- (C) A committee to approve every deck

3. The bad-firm habit of conflating ADOPTION with DISCERNMENT means...

- (A) Everyone uses the tool AND checks it carefully
- (B) 'Every consultant is using the tool' is treated as proof of success — really just proof of a login
- (C) Nobody uses the tool

Ask the teacher — notes to bring to class



4. The unexpected second-order strength of the copilot is...
- (A) Writing recommendations
 - (B) Compressing long transcripts into what the client actually cares about
 - (C) Replacing partners
5. The trap nobody talks about in the town hall is...
- (A) Cost
 - (B) What the productivity gains mean for the pyramid of junior hiring
 - (C) Which vendor to choose

Vocabulary in context

Match each AI-at-work term to its plain-English meaning.

copilot — AI drafting tool integrated with a firm's knowledge base

human-in-the-loop — a named person must review and sign before output goes out

hallucination — confidently-worded but factually wrong AI output

adoption vs. discernment — using the tool ≠ using it well

the pyramid — the traditional junior-heavy staffing model in consulting

review log — a shared record of who signed off on which piece of AI output, when

Vocabulary: AI at work — tools, risks, governance

copilot / assistant — an AI tool that co-authors output alongside a human

prompt / prompting — the instruction you give the model — vague in, vague out

hallucination — confidently wrong output; the classic AI failure mode

grounding (in a source) — forcing the model to cite specific documents

guardrail — a rule the model can't cross (topics, tone, data)

context window — how much text the model can hold in mind at once

human-in-the-loop — a named person MUST scrutinise output before it is dispatched

sign-off / review log — who checked what, when — auditable

red-team (v.) — deliberately try to make the tool fail before a client does

PII (personally identifiable info) — data the tool must NEVER receive

audit trail — the log that lets you reconstruct what happened later

kill switch — one-click way to turn the tool off for a client / team

adoption — how many people log in — a superficial metric in isolation

Ask the teacher — notes to bring to class



discernment — how well people judge what to keep vs. rewrite

productivity uplift — measurable time saved on a defined task

the pyramid — the traditional junior-heavy staffing model

de-skilling risk — juniors never learning the underlying craft

client-facing correction — an error caught after the deliverable went out

practice

Complete each sentence with ONE word / phrase from the sets above.

1. The memo cited a market-share figure that doesn't exist in any source — a textbook ____.

2. No copilot output leaves the firm until a named manager has signed the ____ ____.

3. 60% of our juniors have logged in this quarter. That's ____, not results.

4. Before we ship this to a client, let's ____ - ____ it — try to make it fail on purpose.

5. The tool must never receive ____ from a client transcript — full stop.

6. We measured a 41% ____ ____ on deck drafting — 12 hours down to 7.

Practice 2 — collocations & rewrite

Rewrite each sentence replacing the underlined phrase with ONE word/phrase from this lesson's vocabulary. Keep the meaning identical.

1. The model produced a <u>confidently worded but factually wrong output</u> about our Q3 revenue.

2. We require a <u>named person to scrutinise the output</u> before it goes to the client.

3. We should <u>deliberately try to make the tool fail</u> to find its weak spots.

4. The partners are confusing <u>high usage</u> with actual <u>quality of judgment</u>.

Ask the teacher — notes to bring to class



5. There is a <u>danger that juniors will never learn the craft</u> if they rely too much on AI.

Practice 3 — discursive writing (business memo)

prompt: Write a brief analysis for the firm's leadership regarding the adoption of AI copilots and the necessary governance controls.

Use at least FOUR of the target expressions from this unit.

Adopt a professional register: concise, structured, executive-friendly.

Include a clear recommendation or takeaway in the final sentence.

sample: While we have seen a significant productivity uplift since the introduction of the AI copilot, we must distinguish between simple adoption and genuine discernment. Every consultant is now using the tool, but not all are catching the occasional hallucination that can compromise client deliverables. To mitigate this de-skilling risk, we must enforce a strict human-in-the-loop policy, requiring a manager-grade sign-off on every AI-augmented deck. We also need to implement a centralized review log to track these approvals and provide an audit trail for quality assurance. Red-teaming our internal models has shown that grounding the output in our proprietary research significantly reduces errors, though some risk remains. We must also be vigilant about PII reaching the tool during transcript processing. Recommendation: we should move away from measuring 'usage hours' and instead focus on 'client-facing corrections' to ensure our quality remains uncompromised.

Practice B — AI-at-work collocations

Match each verb to the noun it most naturally goes with.

prompt — the model with a clear instruction and one source

ground — the output in the firm's own documents

review — line-by-line before it touches a client

log — the sign-off with initials and a date

red-team — the tool before the client does

publish — the metric even when it moves the wrong way

Practice C — Build an AI-governance line

Use each cluster of words to write ONE sentence for an internal policy.

1. no copilot output · client-facing · manager-or-above · signed

2. adoption · discernment · we do not measure · alone

3. PII · never · client transcripts · redacted before

Ask the teacher — notes to bring to class



4. hallucinations · monthly · client-facing corrections · publish

Practice D — Odd one out

Which term in each group does NOT belong to the same family?

1. copilot · prompt · grounding · pyramid

2. human-in-the-loop · sign-off · review log · adoption

3. hallucination · guardrail · red-team · kill switch

4. discernment · de-skilling risk · client-facing correction · context window



Grammar focus: quantifiers with precision

1. Universal quantifiers — ALL, EVERY, EACH — ALL + plural noun / uncountable: 'ALL copilot output', 'ALL the juniors'. EVERY + singular countable noun: 'EVERY memo', 'EVERY consultant'. NOT 'every memos' ✗. EACH + singular noun = one at a time, individual view: 'EACH deck is signed by a named manager.' Rule: if there is ONE exception, 'all/every' is FALSE. Downgrade to 'most' or 'almost every'.

2. 'Both' — for exactly TWO — BOTH + plural noun: 'BOTH partners signed.' BOTH ... AND ...: 'BOTH legal AND risk signed off.' Use for exactly TWO items. For three or more, use ALL: 'ALL FOUR partners signed.' 'Both' implies UNANIMITY between the two — check the meaning before you use it in a report.

Ask the teacher — notes to bring to class



3. Large / medium / small quantities — MOST (of the) + plural / uncountable = the majority, more than half. 'MOST junior consultants use it weekly.' MANY + plural = a large number, less specific than 'most'. SOME = an unspecified quantity, more than a couple. A FEW = a small but positive number ('a few clients asked'). FEW (no 'a') = surprisingly small, usually negative in tone ('few consultants read the policy'). HARDLY ANY = almost none — negative framing. NONE / NO = zero. 'NO PII touches the tool.' / 'NONE of the drafts went out unreviewed.'

4. Position and agreement — MOST OF + THE / MY / OUR + noun: 'MOST OF the errors were caught.' ✓ MOST + noun (no 'the') = in general: 'MOST firms have hit their first error.' ✓ 'MOST of firms' ✗ — either 'most firms' OR 'most of THE firms'. With EVERY/EACH the verb is SINGULAR: 'Every consultant HAS a login.' NOT 'Every consultants have...' ✗. With NONE the verb can be singular or plural — singular is safer in writing: 'None of the drafts was sent unreviewed.'

5. Common traps — • 'Every / all people' ✗ → 'Everyone' or 'all people' is fine; 'every people' is not. • 'Most of consultants' ✗ → 'Most consultants' OR 'Most of THE consultants'. • 'Both of them are signed' ≠ 'They are both signed' — same meaning, but the second is what natives write. • 'A few / few' — swapping these reverses your meaning. 'A few clients asked' = positive; 'Few clients asked' = disappointingly few. • 'No / none' — 'no' + noun ('no PII'); 'none' stands alone ('None reached the client').

Exercise A — Choose the precise quantifier

For each sentence, choose the quantifier that is honest, not just impressive.

1. ___ our client-facing decks go through the copilot first — with one exception, a pharma engagement that forbids it. (ALL / MOST)

2. ___ pushback is coming from ___ partners. (MOST ... A FEW / MOST ... FEW)

3. ___ legal ___ risk have signed off. (BOTH ... AND / EVERY ... AND)

4. ___ consultant has completed the AI-use training. (EVERY / ALL)

5. ___ PII from client transcripts reaches the tool. (NO / NONE)

Exercise B — Rewrite the exaggeration

Each line is too strong to be true. Rewrite it with a HONEST quantifier.

1. 'Every consultant is already using the tool.'

2. 'All our decks are reviewed line-by-line.'

Ask the teacher — notes to bring to class



3. 'No consultant would ever send an AI draft without checking.'

4. 'Both risk and legal love it.'

Exercise C — Fix the grammar

Correct the quantifier error in each sentence.

1. 'Every consultants have completed the training.'

2. 'Most of firms have hit their first client-facing error.'

3. 'A few clients asked — hardly a signal.'

4. 'No PII touch the tool.'

Listening: three partners on the copilot rollout

You're the Managing Partner. Three partners have 90 seconds each to give you their read on the firm-wide copilot rollout. You're listening for: (1) who uses quantifiers HONESTLY, (2) who conflates ADOPTION with DISCERNMENT, (3) who names the ONE thing that would change their mind, (4) whose recommendation would survive contact with a Tier-1 client.

pre

Which partner will you trust most before you hear them? And is that because of what they know, or because of how they sound?

What ONE quantifier ('all', 'most', 'a few') will you flag every time it's used loosely?

gist

1. Which partner's recommendation gets adopted, and why?

2. Which partner runs on ADOPTION dressed as success?

3. Which partner has good instincts but no discipline?

Ask the teacher — notes to bring to class



detail

1. What EXACTLY is Anika's rule for a draft that touches a client?

2. What is Julian's answer when asked about the one client-facing error last month?

Third listen — designer vs default

Focus on the *QUANTIFIERS* each partner reaches for. Answer from what they say, not from what they seem to mean.

1. Anika uses 'MOST ... with a carve-out', 'A FEW clients', 'HARDLY ANY'. What is that showing you as MP?

2. Julian uses 'EVERY junior', 'ALL our decks', 'MASSIVE uplift'. What does that language reveal?

3. Diego says 'in most cases', 'sort of', 'we should probably'. What is the risk you carry away?

Language from the audio — quantifier moves worth stealing

Match each phrase to what it *DOES* for the speaker.

"MOST client-facing decks — with a flagged carve-out for the pharma engagement." — downgrades 'all' to 'most' AND names the exception — the honest partner move

"A FEW Tier-1 clients have proactively asked — I read that as the first three of many." — hedged minority upgraded to a signal — reads the leading edge, doesn't wait for 'most'

"HARDLY ANY drafts went out unreviewed — the ones that did are in the log." — admits the number is not zero AND shows the audit trail exists — trust through evidence, not claim

"EVERY junior is using it — MASSIVE uplift across the board." — universal claim + intensifier + no numbers = adoption dressed as success; the tell of a partner who hasn't audited

"In MOST cases, we sort of review the output — most of the time." — hedged frequency with no rule and no signer — good instinct, zero discipline

Ask the teacher — notes to bring to class



"NO copilot output leaves this firm to a client until a named manager has signed the log." — the discipline in one sentence — 'no' + named signer + auditable log

Voice A (Julian): Look, the numbers speak for themselves. EVERY junior on my team is using the copilot — I mean literally every one. ALL our client-facing decks go through the tool first. We're seeing MASSIVE uplift across the board — drafting time is down, morale is up. Yes, there was that one memo last month, but honestly, one-off. Everyone learned from it. We should be rolling this out to all 400 consultants yesterday.

Voice B (Anika): OK. Here's what I actually see in my practice. MOST client-facing decks — with a clearly-flagged carve-out for the pharma engagement, which the contract forbids — go through the copilot first. A FEW Tier-1 clients have proactively asked whether we use AI on their work, and I read that as the first three of many. HARDLY ANY drafts went out unreviewed last quarter, but the ones that did are in the log — three of them, all in the same week, one signer over-loaded. My rule is simple: NO copilot output leaves this firm to a client until a named manager at grade above has signed the log with initials and a date. Two client-facing corrections in a quarter and I pull the tool from that practice — that's the number that would change my mind.

Voice C (Diego): I mean, I hear both of them. In MOST cases, we sort of review the output — most of the time. I think we should probably have some kind of sign-off. The risk is real, obviously — this can't just be a login, right? But I don't want to slow the team down with too much process. Maybe a light-touch review? We should probably discuss it more.

Speaking: solo role-play — 90-second copilot brief to your MP

Solo task. You are one of the three partners. Give your 90-second brief to the Managing Partner. Record yourself (voice note, up to 90 seconds). You MUST: (1) use at least THREE quantifiers HONESTLY (name at least one exception when you say 'most' or 'all'); (2) name your ONE human-in-the-loop RULE in a single testable sentence; (3) name the ONE metric that would change your mind — and in which direction; (4) end with what YOUR practice will publish monthly.

Scene: Your real current team (recommended)

A: You brief.

B: Imagined listener: the Managing Partner, 90 seconds, no slides.

Scene: A firm you're advising as a consultant

A: You brief.

B: Imagined listener: their MP, same rules, same time.

Ask the teacher — notes to bring to class



Scene: A department that has just piloted an AI tool

A: You brief.

B: Imagined listener: the CEO who signed the pilot budget.

Extra speaking task — 30-second 'one rule, one metric'

Compress your brief to **ONE** rule you will **INSTALL** next week — and the **ONE** metric that would kill the rollout. Record a **30-second voice note**. Format: 'Starting Monday, no ____ until _____. If _____ in a quarter, we pull the tool from that practice.'

Writing mini-task

Write the **ONE**-page 'How this firm uses AI' policy (max 130 words). It must contain: (1) the **HUMAN-IN-THE-LOOP** rule in one sentence; (2) the **THREE** monthly **METRICS** (one speed, one quality, one trust); (3) at least **ONE** quantifier ('most / a few / hardly any') used **HONESTLY**, with a named exception. Sign it with a real role.

Wrap-up

One **QUANTIFIER** from today (most / a few / hardly any / none) you'll use with a **NAMED EXCEPTION** next time.

One universal claim ('every...', 'all...') you'll actively **STOP** using without a carve-out.

One line for your team's AI-use policy that you'll draft this week and get signed by Friday.

Quiz

- Which sentence uses quantifiers **HONESTLY**?
 - 'Every consultant is using the tool and every deck is perfect.'
 - 'Most client-facing decks go through the tool — with a flagged carve-out for the pharma engagement.'
 - 'All feedback has been massively positive.'
- The correct pattern is...
 - 'Every consultants have completed the training.'
 - 'Every consultant has completed the training.'
 - 'Every consultants has completed the training.'
- 'Both' is correctly used when...
 - Three teams have signed off
 - Exactly two teams / people have signed off
 - You want to sound emphatic

Ask the teacher — notes to bring to class



4. The safest 'human-in-the-loop' rule reads...
- (A) 'We try to review most of the output.'
 - (B) 'No copilot output reaches a client until a named manager-or-above has signed the review log (initials + date).'
 - (C) 'Legal has approved the tool.'
5. 'Most of firms have shipped their first error' is...
- (A) Correct
 - (B) Wrong — 'MOST FIRMS' or 'MOST OF THE FIRMS', never 'most of firms'
 - (C) Correct in informal writing

Ask the teacher — notes to bring to class



Workbook — home practice



Exercise 1 — AI-at-work vocabulary

Complete each sentence with *ONE* phrase from the box.

Word bank: hallucination · human-in-the-loop · review log · red-team · adoption · discernment · kill switch

1. The memo cited a market-share figure with no basis in any source — a textbook ____.

2. No copilot output goes to a client until a ____ - ____ - ____ - ____ has signed the log.

3. Every sign-off is captured in the ____ ____ — initials, date, link.

4. Before we ship this externally, let's ____ - ____ it — try to make it fail on purpose.

5. 70% of juniors logged in this quarter — that constitutes ____, not results.

6. The number we can't measure yet is ____ — how well people judge what to keep vs. rewrite.

7. If the client-facing correction rate doubles, we hit the ____ ____ and turn it off for that account.

Ask the teacher — notes to bring to class



Exercise 1b — Vocabulary — discursive writing

Write a policy proposal (120–180 words) for the Partners regarding the implementation of Generative AI in the firm. Incorporate at least four phrases from Exercise 1: *hallucination*, *human-in-the-loop*, *review log*, *red-team*, *adoption*, *discernment*, or *kill switch*.

type: writing

Use at least 4 target vocabulary items from the unit.

Maintain a formal, cautious yet forward-looking register.

Use precise quantifiers (e.g., 'most', 'hardly any', 'a few').

Structure clearly: Opportunity, Risk Mitigation, Performance Metrics.

model: Subject: Policy for Generative AI Implementation To the Partners, The pilot results for our generative AI integration show high levels of adoption, yet we must remain vigilant regarding output quality. We have identified several instances of hallucination where the tool fabricated market data, reinforcing the necessity of our 'human-in-the-loop' requirement. No AI-generated deliverable may be sent to a client without being recorded in the review log by a manager. To further harden our defences, I recommend we red-team our most sensitive prompts to identify potential failure modes. Furthermore, we must invest in training to improve the discernment of our staff—ensuring they can distinguish between high-quality insights and generic filler. If the rate of uncorrected errors exceeds 2%, we will activate the kill switch for that specific use case. Our goal is to leverage efficiency without compromising client trust.

Exercise 2 — Quantifier precision

Rewrite each exaggerated sentence with an HONEST quantifier ('most', 'a few', 'hardly any', 'none') and — where relevant — name the exception.

1. 'Every consultant is already using the tool.'

2. 'All our client-facing decks go through the copilot first.'

3. 'No consultant would send an AI draft without checking.'

4. 'Both risk and legal love it.'

5. 'A few Tier-1 clients have asked — hardly a signal.'

Exercise 3 — Quantifier grammar

Correct the ONE error in each sentence.

Ask the teacher — notes to bring to class



1. 'Every consultants have completed the training.'

2. 'Most of firms have hit their first client-facing error.'

3. 'Both of the four partners signed off.'

4. 'No PII touch the tool.'

5. 'None of the drafts were sent unreviewed, everyone was.'

Exercise 4 — Honest or adoption-dressed-as-success?

Decide whether each line is *HONEST* (uses quantifiers with a named exception / consequence) or *ADOPTION-AS-SUCCESS* (universal, unnumbered, unfalsifiable).

1. 'Most client-facing decks — with a flagged carve-out for the pharma engagement — go through the copilot.'
(A) HONEST
(B) ADOPTION-AS-SUCCESS
2. 'Every junior is using the tool — massive uplift across the board.'
(A) HONEST
(B) ADOPTION-AS-SUCCESS
3. 'Hardly any drafts went out unreviewed — the ones that did are in the log.'
(A) HONEST
(B) ADOPTION-AS-SUCCESS
4. 'All feedback has been positive.'
(A) HONEST
(B) ADOPTION-AS-SUCCESS
5. 'A few Tier-1 clients have asked — I read it as the first three of many.'
(A) HONEST
(B) ADOPTION-AS-SUCCESS
6. 'Everyone learned from the one client-facing error last month.'
(A) HONEST
(B) ADOPTION-AS-SUCCESS

Ask the teacher — notes to bring to class



Exercise 5 — Rewrite the hype line as a policy line

Rewrite each hype line as a testable one-sentence policy that a manager would sign under.

1. 'We're rolling AI out firm-wide — the future is here.'

2. 'AI will save junior consultants massive amounts of time.'

3. 'Every consultant will be trained on the tool.'

4. 'Client data is safe.'

Exercise 6 — Three partners on the copilot rollout

You'll hear three partners brief the Managing Partner on their practice's copilot rollout. Decide which recommendation you'd sign under.

1. Whose recommendation would you take to the partner call, and why?

2. Whose brief is adoption dressed as success?

3. Whose brief has instinct but no discipline?

Voice A (Julian): Every junior on my team is using the copilot — literally every one. All our client-facing decks go through the tool first. Massive uplift across the board. Yes, there was that one memo last month, but it's a one-off. Everyone learned from it. Roll to 400.

Voice B (Anika): MOST client-facing decks — with a flagged carve-out for the pharma engagement — go through the copilot. A FEW Tier-1 clients have proactively asked; I read it as the leading edge. HARDLY ANY drafts went out unreviewed; the exceptions are in the log. My rule: NO copilot output leaves this firm to a client until a named manager-or-above has signed the log. Two client-facing corrections in a quarter and I pull the tool from that practice — that's the number.

Voice C (Diego): In most cases we sort of review the output — most of the time. I think we should probably have some kind of sign-off. The risk is real, this can't just be a login. But I don't want to slow the team down. Maybe a light-touch review. We should probably discuss it more.

Ask the teacher — notes to bring to class



Exercise 7 — Write your team's AI-use policy

Write your OWN team's ONE-PAGE 'How we use AI' policy (max 130 words). Cover: (1) the human-in-the-loop rule, in one testable sentence; (2) the three monthly metrics (speed, quality, trust); (3) at least ONE quantifier ('most', 'a few', 'hardly any') used HONESTLY with a named exception. Sign it with a real role.

Ask the teacher — notes to bring to class



Answer key — Class Book

Lead-in: name the quantifier trap

1. → UNIVERSAL claim, almost never true — the loud minority stands in for the median.
2. → CONTRADICTION inside one sentence — 'most' pushback from 'few' people can't both be true.
3. → FALSE TOTALITY — one exception ('this pharma client won't allow it') and the claim collapses.
4. → PAIRED-UP SHIELD — 'both' implies unanimity; sign-off means 'not illegal', not 'good idea'.
5. → ABSOLUTE DENIAL — every 'no' at this scale hides at least one Tuesday afternoon at 17:55.
6. → HEDGED MINORITY — 'a few' is soft cover for 'this is now on every Tier-1 client's radar and we should get ahead of it.'

Reading: 'It drafts. We decide.'

1. (B) Which firms build the discipline to keep the time savings without losing the client
2. (B) Line-by-line review by a named manager-or-above who signs a shared log
3. (B) 'Every consultant is using the tool' is treated as proof of success — really just proof of a login
4. (B) Compressing long transcripts into what the client actually cares about
5. (B) What the productivity gains mean for the pyramid of junior hiring

Vocabulary: AI at work — tools, risks, governance

1. hallucination
 2. review log
 3. adoption
 4. red-team
 5. PII
 6. productivity uplift
1. The model produced a hallucination about our Q3 revenue.
 2. We require a human-in-the-loop before it goes to the client.
 3. We should red-team the tool to find its weak spots.
 4. The partners are confusing adoption with actual discernment.
 5. There is a de-skilling risk if juniors rely too much on AI.
1. No copilot output goes to a client until a manager-or-above has signed the review log — no exceptions.
 2. Adoption is not a KPI on its own — we measure discernment (kept vs. rewritten) alongside it.
 3. PII from client transcripts is redacted before any text reaches the tool — never after.
 4. We publish client-facing corrections per 100 deliverables monthly, so hallucinations show up as a number, not a rumour.
1. pyramid — the other three are about USING the tool; pyramid is about staffing structure.
 2. adoption — the other three are governance controls; adoption is a usage metric.
 3. hallucination — the other three are DEFENCES; hallucination is the FAILURE they defend against.
 4. context window — the other three are consequences of how humans use the tool; context window is a technical limit of the model.

Grammar focus: quantifiers with precision

1. MOST — 'all' is false because of the pharma exception. 'Most' with an explicit carve-out is the honest word.
 2. MOST ... A FEW — but flag the contradiction to the speaker. If it's really 'a few' partners, it can't really be 'most' pushback.
 3. BOTH ... AND — 'both' is right because there are exactly TWO teams.
 4. EVERY — 'every' takes a SINGULAR noun ('consultant'). 'All consultants have...' would also be right, but the two are not interchangeable grammatically.
 5. NO — 'no' takes a noun. 'NONE reaches the tool' would also work ('none' stands alone).
1. 'MOST junior consultants use the tool weekly; a smaller share of partners have adopted it.'
 2. 'The MAJORITY of client-facing decks are reviewed line-by-line; the exceptions are logged and reviewed at the following week's quality huddle.'
 3. 'HARDLY ANY consultant sends a draft without checking — but the audit trail exists precisely because the exception is not zero.'

Ask the teacher — notes to bring to class



4. 'BOTH risk AND legal have signed off that the tool is contractually usable; neither has yet endorsed unrestricted use.'
1. 'EVERY CONSULTANT HAS completed the training.' (Every + singular noun + singular verb.)
2. 'MOST FIRMS have hit their first error.' OR 'MOST OF THE FIRMS have hit their first error.' (Not 'most of firms'.)
3. Depending on meaning: 'A FEW clients asked' = 'some did, it's real' (positive). If the speaker means 'not many', use 'FEW clients asked — hardly a signal.'
4. 'NO PII TOUCHES the tool.' ('PII' is uncountable / treated as singular; use a singular verb.)

Listening: three partners on the copilot rollout

1. Voice B (Anika). She uses quantifiers honestly: 'MOST decks — with a clearly-flagged carve-out for the pharma engagement', 'A FEW Tier-1 clients have proactively asked', 'HARDLY ANY drafts went out unreviewed — the ones that did are in the log'. Names the ONE thing that would change her mind: two client-facing corrections in a quarter.
2. Voice A (Julian). 'EVERY junior is using it', 'ALL our decks now go through the tool', 'we're seeing MASSIVE uplift'. No numbers, no exceptions, no named person doing the review — logins as proof. This is exactly the language that produces a Tier-1 client catching the error before you do.
3. Voice C (Diego). Sees the risk ('this can't just be a login'), but names no rule, no signer, no metric. Everything is 'we should probably', 'kind of', 'in most cases'. Six weeks from now, his practice will have both Anika's caveats AND Julian's blind spots — the worst of both.
 1. No copilot output — memo, slide, model — leaves the firm to a client, or is presented in a client-facing meeting, until a named consultant at Manager grade or above has signed the review log (initials + date) confirming (a) every number checked against source and (b) the recommendation is one they would defend in the room.
 2. He calls it 'a one-off', says 'EVERYONE learned from it', and moves back to the 41% deck-drafting uplift. He never names WHO reviewed the memo before it went out, or what specifically changed in the process after the client caught it. That's the tell.
 1. She is telling you where the exceptions are BEFORE you find them yourself. That is what a partner sounds like when they've genuinely walked their own book — the number is downgraded from 'all' to 'most' because she can name the one client that isn't.
 2. It reveals that he is presenting adoption as evidence of quality. 'Every' at 400-consultant scale is almost never true — one Tuesday afternoon at 17:55 breaks it. The strong quantifier is the tell: he has not gone looking for the exception, which means the client will.
 3. He has correctly diagnosed the risk (adoption ≠ discernment) but he has no RULE and no SIGNER. Hedged instincts do not survive contact with a busy Tuesday. Within a quarter, his practice will look exactly like Julian's — same output, same first client-facing error, but with the added trap that Diego 'saw it coming' and didn't act.

Quiz

1. (B) 'Most client-facing decks go through the tool — with a flagged carve-out for the pharma engagement.'
2. (B) 'Every consultant has completed the training.'
3. (B) Exactly two teams / people have signed off
4. (B) 'No copilot output reaches a client until a named manager-or-above has signed the review log (initials + date).'
5. (B) Wrong — 'MOST FIRMS' or 'MOST OF THE FIRMS', never 'most of firms'

Ask the teacher — notes to bring to class



Answer key — Workbook

Exercise 1 — AI-at-work vocabulary

1. hallucination
2. human-in-the-loop
3. review log
4. red-team
5. adoption
6. discernment
7. kill switch

Exercise 2 — Quantifier precision

1. 'MOST junior consultants use the tool weekly; adoption among senior partners is significantly lower.'
2. 'MOST client-facing decks go through the copilot first — with a flagged carve-out for the pharma engagement, which the contract forbids.'
3. 'HARDLY ANY consultant sends a draft without checking — the exceptions are in the review log, three cases last quarter.'
4. 'BOTH risk AND legal have signed off that the tool is contractually usable; neither has endorsed unrestricted use.'
5. 'A FEW Tier-1 clients have proactively asked in the last month; I read that as the leading edge, not the whole wave.'

Exercise 3 — Quantifier grammar

1. 'EVERY consultant HAS completed the training.' (Every + singular noun + singular verb.)
2. 'MOST FIRMS have hit their first error.' OR 'MOST OF THE FIRMS have hit their first error.'
3. 'ALL FOUR partners signed off.' (BOTH is for exactly TWO.)
4. 'NO PII TOUCHES the tool.' (PII treated as uncountable → singular verb.)
5. 'NONE OF THE DRAFTS WAS sent unreviewed.' (In writing, NONE + singular verb is safer and more consulting-register.)

Exercise 4 — Honest or adoption-dressed-as-success?

1. (A) HONEST
2. (B) ADOPTION-AS-SUCCESS
3. (A) HONEST
4. (B) ADOPTION-AS-SUCCESS
5. (A) HONEST
6. (B) ADOPTION-AS-SUCCESS

Exercise 5 — Rewrite the hype line as a policy line

1. 'From September 1st, the copilot is available to 90 consultants across four offices. It drafts. We decide. No copilot output goes to a client until a Manager-or-above has signed the review log — no exceptions.'
2. 'We will publish monthly: (1) median hours from brief to first-draft deck, (2) client-facing corrections per 100 deliverables, (3) % of Tier-1 clients proactively told about the tool and NOT opting out.'
3. 'Every consultant using the tool must complete a 45-minute AI-use module and sign the acceptable-use policy before their login is enabled — logged and dated in HR.'
4. 'PII from client transcripts is redacted BEFORE any text reaches the tool. Redaction is verified by a named consultant, initialled in the log. Two failures in a quarter and the practice loses access.'

Exercise 6 — Three partners on the copilot rollout

1. Voice B (Anika). Uses quantifiers honestly ('MOST with a carve-out', 'A FEW clients', 'HARDLY ANY'), names the review rule ('no output leaves until a named manager signs the log'), names the kill-metric (two client-facing corrections in a quarter → tool pulled).
2. Voice A (Julian). 'EVERY junior', 'ALL decks', 'MASSIVE uplift' — universal claims, no named signer, no kill-metric, deflects on the one client-facing error.

Ask the teacher — notes to bring to class



3. Voice C (Diego). Correctly diagnoses the risk ('this can't just be a login') but hedges every sentence ('in most cases', 'sort of', 'we should probably') and names no rule, no signer, no metric.

Ask the teacher — notes to bring to class
